Nested model F tests activity: solutions

STAT3530 Applied Linear Models

Solutions to the blanks in `scripts/10-f-tests.R`.

```
seat <- lm(hipcenter ~ ., data = seatpos)
```

1. Testing a subset

Each test fits the reduced model H_0 describes and compares it to the full model.

```
# (a) do the body measurements matter, given Age and Weight?
anova(lm(hipcenter ~ Age + Weight, data = seatpos), seat)
```

```
Analysis of Variance Table

Model 1: hipcenter ~ Age + Weight
Model 2: hipcenter ~ Age + Weight + HtShoes + Ht + Seated + Arm + Thigh +
    Leg
  Res.Df   RSS Df Sum of Sq      F  Pr(>F)
1     35 68923
2     29 41262  6     27661 3.2402 0.01465 *
---
Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
```

```
# (b) do we need both HtShoes and Ht, given everything else?
anova(lm(hipcenter ~ . - HtShoes - Ht, data = seatpos), seat)
```

```
Analysis of Variance Table

Model 1: hipcenter ~ (Age + Weight + HtShoes + Ht + Seated + Arm + Thigh +
    Leg) - HtShoes - Ht
Model 2: hipcenter ~ Age + Weight + HtShoes + Ht + Seated + Arm + Thigh +
    Leg
  Res.Df   RSS Df Sum of Sq      F Pr(>F)
1     31 42239
2     29 41262  2    976.85 0.3433 0.7123
```

```r
# (c) does Age matter, given everything else?
anova(lm(hipcenter ~ . - Age, data = seatpos), seat)
```

```
Analysis of Variance Table

Model 1: hipcenter ~ (Age + Weight + HtShoes + Ht + Seated + Arm + Thigh +
    Leg) - Age
Model 2: hipcenter ~ Age + Weight + HtShoes + Ht + Seated + Arm + Thigh +
    Leg
  Res.Df   RSS Df Sum of Sq      F Pr(>F)
1     30 43894
2     29 41262  1    2632.1 1.8499 0.1843
```

How can a group matter when no single one does? The body measurements are nearly collinear, so each is redundant *given the others* — dropping any one costs almost nothing. Dropping all six at once costs a lot, because then nothing is left to stand in for them.

```r
round(summary(seat)$coefficients[, 3:4], 4)
```

```
            t value Pr(>|t|)
(Intercept)  2.6201   0.0138
Age          1.3601   0.1843
Weight       0.0795   0.9372
HtShoes     -0.2761   0.7845
Ht           0.0594   0.9531
Seated       0.1419   0.8882
Arm         -0.3405   0.7359
Thigh       -0.4297   0.6706
Leg         -1.3660   0.1824
```

2. One predictor, two ways

```r
summary(seat)$coefficients['Age', 't value']^2
```

```
[1] 1.84993
```

Identical to F in (c): a one-constraint F test *is* the partial t test, squared.

3. Testing a contrast

The constraint $\beta_{\texttt{HtShoes}} = \beta_{\texttt{Ht}}$ forces the two to share a coefficient, so they enter the reduced model as a sum.

```
byhand <- lm(hipcenter ~ Age + Weight + I(HtShoes + Ht) + Seated + Arm +
               Thigh + Leg, data = seatpos)

anova(byhand, seat)
```

```
Analysis of Variance Table

Model 1: hipcenter ~ Age + Weight + I(HtShoes + Ht) + Seated + Arm + Thigh +
    Leg
Model 2: hipcenter ~ Age + Weight + HtShoes + Ht + Seated + Arm + Thigh +
    Leg
  Res.Df   RSS Df Sum of Sq      F Pr(>F)
1     30 41302
2     29 41262  1    39.755 0.0279 0.8684
```

Same $F = 0.0279$ and $p = 0.8684$ as `car::linearHypothesis`.

Does this tell you whether height matters? No. It says the two height measures can share a coefficient, not that either one does anything — (b) already showed you can drop both.

Your turn: savings rates

```
m1 <- lm(sr ~ ., data = savings)
```

A subset. $H_0 : \beta_{\texttt{pop15}} = \beta_{\texttt{pop75}} = 0$.

```
anova(lm(sr ~ dpi + ddpi, data = savings), m1)
```

```
Analysis of Variance Table

Model 1: sr ~ dpi + ddpi
Model 2: sr ~ pop15 + pop75 + dpi + ddpi
  Res.Df    RSS Df Sum of Sq      F   Pr(>F)
1     47 824.72
2     45 650.71  2       174 6.0167 0.004835 **
---
Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
```

Age structure matters: $F = 6.02$ on 2 and 45 df, $p = 0.005$.

One predictor.

```
anova(lm(sr ~ . - ddpi, data = savings), m1)
```

```
Analysis of Variance Table

Model 1: sr ~ (pop15 + pop75 + dpi + ddpi) - ddpi
Model 2: sr ~ pop15 + pop75 + dpi + ddpi
  Res.Df    RSS Df Sum of Sq      F  Pr(>F)
1     46 713.77
2     45 650.71  1    63.054 4.3605 0.04247 *
---
Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
```

```
summary(m1)$coefficients['ddpi', 't value']^2
```

```
[1] 4.360496
```

$F = t^2 = 4.36$, $p = 0.042$, as before.

A contrast.

```
car::linearHypothesis(m1, 'pop15 + pop75 = 0')$`Pr(>F)`[2]
```

```
[1] 0.07904611
```

```
car::linearHypothesis(m1, 'pop15 = pop75')$`Pr(>F)`[2]
```

```
[1] 0.2145833
```

Different questions: the first asks whether the two effects cancel, the second whether they are the same number. Neither is rejected here, but that is not evidence the demographic effects are absent — the subset test above already rejected $\beta_{\texttt{pop15}} = \beta_{\texttt{pop75}} = 0$.